Review of CAS Development and Deployment Effort

Fermi National Accelerator Laboratory, March 1-2, 2004

Report of the Review Committee

Amber Boehnlein, Liz Buckley-Geer, Daniel Durand, Lee Hurt,
Scott Kleinman, Lee Lueking, Jeff Munn

March 12, 2004

1 Executive Summary

The CAS provides an exceptional public interface to the SDSS data. Its design is well thought out, providing simple and efficient interfaces for the public and professional astronomers alike, as well as a powerful SQL-based interface for the sophisticated user, capable of taking full advantage of the complex data model and efficient database design. It contributes significantly to the project's public outreach commitment. The range of interfaces and functionality is truly impressive, and exceeds the requirements.

While the functionality is exceptional, the CAS is not yet a fully operational system. The project's top level plan is for a turn-key system based at FNAL, fully integrated with the current Data Processing system. This is important both to allow the timely dissemination of data to the collaboration and public, as well as for long term support issues. The committee agrees with this top level plan.

The lack of agreed upon operational requirements, and an implementation plan based on those requirements, has been the primary impediment to delivery of a fully integrated operational CAS. There is a general sense that the JHU and EAG teams have often disagreed on priorities, and that development has proceeded based on inadequate communication and agreement between the two teams. While adequate functionality requirements are in place, there are no operational requirements. Even were there operational requirements, the committee feels that requirements are not driving development. The result has been that functionality has generally taken precedence over operations. This has contributed to delays in data releases, both to the collaboration and the public (though there have been equally important causes of delay upstream of

the CAS). A detailed implementation plan, including schedule and manpower estimates and clear milestones for all aspects of the deployment of an integrated operational CAS, was not presented.

While there are no technical concerns over the hardware plan, it lacks sufficient detail concerning such issues as manpower, support (both short-term and long-term), and its relation to FNAL as a whole. The technical aspects of choice of CPUs, disk drives, etc raises no concerns for the committee, beyond a certain sense of "fragility" due to financial constraints. However, there are not an adequate set of performance requirements by which to evaluate the plan. There is concern that the support of a suite of Windows machines running Microsoft SQL Server, including the database administration, has not been adequately planned. The EAG group has neither the Windows nor SQL Server expertise, nor the manpower, to adequately support all aspects of the CAS autonomously from the JHU group. The committee commends the collaboration with the CSS/CSI group, and hopes that by leveraging the Windows expertise within that group, sufficient knowledge can be obtained to establish a reliable and usable system.

There are no requirements on the CAS concerning collaboration access. To date, the collaboration has been offered minimal proprietary access to SDSS data via the CAS. Collaboration needs have been secondary to public releases. This has been a source of contention within the collaboration, and has limited use of the CAS by the collaboration.

Heroic efforts have been and are being made both at JHU and FNAL to develop and deploy the CAS. The data releases to date have been deployed under tremendous time pressure. Again, the lack of operational requirements, and an agreed upon implementation plan based on the requirements with clear milestones, has contributed to this pressure, as have delays in processing upstream of the CAS. This level of effort cannot be maintained without risking burnout. It is recognized by the committee that DR2 must be deployed before the recommendations in this report can be fully considered.

With the above in mind, the committee makes the following key recommendations:

- Re-evaluate and update the requirements within one month of the DR2 public data release. Particular attention should be paid to the need for operational and performance requirements.
- Write a detailed implementation plan, based on the requirements and agreed upon by both the JHU and EAG groups, and with clear priorities and milestones, within the same time frame.
- Make the top priority, after DR2, the deployment of a robust CAS at FNAL, fully integrated with Data Processing, capable of efficiently and reliably loading the entire DR3 data set autonomously from the JHU group.
- Freeze data model changes until DR4 at the earliest.

- Revisit the hardware plan, with an eye toward manpower, support, integration, networking, security, etc.
- Proceed with purchasing of the machines for DR3, prior to the reassessment of requirements and plans.
- Supply one full-time support person at FNAL for the CAS with Windows and SQL Server expertise.
- Seek continued collaboration with FNAL/CD.
- Write MOUs between EAG and JHU, and EAG and the relevant FNAL divisions/departments, to understand each others roles in the long term support of the CAS.
- Define the role of the CAS for the collaboration, and incorporate that role in the requirements and implementation and hardware plans.

2 The Report

The charge to the review committee was in the form of four questions, each of which is addressed in detail here in the form of findings, comments, and recommendations. Where appropriate, references are made to the July 2000 "SDSS Data Processing and Distribution Review" report, as requested in the charge to the committee.

2.1 Are the requirements for the CAS well-defined? If not, what areas should be clarified?

Findings

- *The requirements that were presented to the committee are out-of-date.* They were initially developed in 1995 and have not been substantially revisited since then. For example they still contain references to object databases which are no longer part of the plan.

Comments

- *The functional requirements are well-defined.* The specific data that must be stored, the types of queries that must be supported, and the user interfaces are all fully constrained by a complete and detailed set of requirements.
- *The performance and operational requirements are not well-defined.* Uptime requirements were added a day or two before the review, however there are no requirements on, for example, the number of concurrent users

or the size of data sets returned to the user. Just as importantly, there are no requirements concerning the operational performance of the database. For example, while there is a strong desire for the CAS to be loaded and operated autonomously from the developers, there are no requirements to guide the development of the CAS in that direction or provide milestones by which progress may be measured. Integration of the data processing pipelines into a smooth operation has benefited greatly from requirements (though possibly not formal ones) on the execution (and failure) modes of the pipelines, execution times, quality analysis, etc; integration of CAS into FNAL operations would benefit from similar requirements.

- *Existing requirements are not driving development.* The committee feels that the current development work is not being driven by the requirements but rather by ad-hoc responses to users requests or singular visions of pertinent directions. The lack of up-to-date requirements also makes it hard to separate work required to provide core SDSS functionality from work related to other areas such as NVO. The implications of each requirement should be carefully thought through; for example the 99% uptime requirement for the public CAS, and the number of Data Releases to be maintained at a given time. There are features which could have large development costs as well as hardware and support implications that are the result of this lack of well-defined performance and operational requirements; one such example is the CasJobs/MyDB product.

- *There are no requirements specifying the role of the CAS in providing data to the collaboration.* The short proprietary time that the collaboration has had access to the CAS prior to public releases, and the fact that all data is not available, has been a source of considerable contention within the collaboration, and driven many potential CAS users to "roll their own". The role of CAS in providing data to the collaboration should be clarified; a set of requirements specifying which data is to be made available via the CAS, and on what time scale, would help both to clarify, and to determine the priority for, CAS's role in providing data to the collaboration. The important role the collaboration provides in verifying both the functionality of the CAS and the data integrity should not be overlooked. This was emphasized in the 2000 review report: "The Science Databases are the means by which the main body of the collaboration has access to the data and is enabled to do quality and scientific analysis. During this first year of production data taking, it is crucial that as many people in the collaboration as possible have access to and work with the acquired data in a timely fashion and provide in depth verification and feedback on the performance of the instruments."

Recommendations

- Re-evaluate and update the requirements within one month after DR2 is completed. The requirements should clearly differentiate between mandatory features and desirable features. Requirements for the CAS versus the SkyServer should be defined separately. The implications of each requirement should be carefully considered, particular regarding hardware, support, and manpower.
- Add performance requirements. These might include the required system uptime, the maximum number of users to be supported, the maximum size of data sets to return to the users, etc.
- Add operational requirements. These might include robustness, integration with FNAL Data Processing, ability for incremental loading, ability to back out chunks, documentation, etc.
- Account for long term support and operations when re-evaluating the requirements.
- Add requirements specifying the role of the CAS in providing data to the collaboration (i.e., which data and on what time scale).

2.2 Is there a clear path to completion of the CAS development?

Findings

- *There is not an agreed upon, prioritized plan for the completion of the CAS development.* There are task lists for both hardware deployment and software feature development that were provided in the review. However, there is not a work list agreed to by both JHU and FNAL/EAG, with a prioritized schedule of tasks tied to milestones, which would coordinate the manpower devoted to the project and help assure its successful completion. The management recognizes this lack of a detailed implementation plan. The need for such a plan was identified in the 2000 review report: "The presentation on data distribution did not include any discussion of how the final data products would be generated, validated, and documented. The complexity and drain on resources for these tasks should not be underestimated. The committee recommends that the project define and scope the job of preparing data for release before plans for either the early or subsequent releases are finalized."

Comments

- *The lack of up-to-date requirements contributes to the imprecise plan.* Many of the features being implemented are not in the requirements, and

the selection and deployment of the hardware is not requirements driven.

- *The lack of a mutually agreed upon plan has slowed the delivery of a robust CAS system.* Frequent, un-planned changes to the system have required additional effort and delayed the delivery of the CAS data releases to the collaboration and public. Delays upstream of the CAS have been equally important in delaying data releases, yet CAS functionality which would aid in dealing with upstream delays, such as incremental loading and the ability to back out changes, have not been prioritized. The committee feels that adding features has taken priority over operations. Indeed, the 2000 review committee had the same concern: "While development of support of more query modes is needed for the Science Databases, there should also be an emphasis on ensuring smooth and timely transfer of data from the Operational to the Science Databases as part of the standard data processing production cycle."

- *No data validation plans were presented.* As such, it is impossible for the committee to address whether adequate means for data validation are in place. This was a particular concern of the 2000 review committee: "The project should also define a set of validation tests and criteria to establish that the data products are of release-quality."

Recommendations

- Generate an implementation plan including details for both software and hardware. This plan should be developed and vetted within one month after DR2 is completed, and in accordance with the updated requirements.

- Tie the implementation of requirements and features to major milestones of the experiment. To assure the system is available in a timely fashion with each data release, functional milestones should be closely related to the DR schedule. Timely releases of the data will strongly enhance the use of the system by the collaboration.

- Establish intermediate milestones and checkpoints within the plan and evaluate progress based on them. In order to meet the major milestones and assure success, intermediate goals are essential.

- Establish a procedure for change control, similar to that used for APO operations. The SDSS collaboration does have a change control board, but it is felt that the timeliness of the model used at APO is more appropriate for this critical area.

- Freeze the development of CAS after DR2 until robust operations are achieved. We concur with management that the efficient, repeatable loading of DR3 at FNAL, autonomously from the JHU group, is the minimal

test of a fully operational CAS. Such a fully operational system must be the top priority after DR2.

- Prohibit any data model changes until DR4 at the earliest. Changes to the data model prior to and during past releases have contributed to long delays and significant additional effort to deploy the CAS at Fermilab. Having a stable schema will significantly improve the situation, allowing development to concentrate on operational issues, and enable the CAS data to be available to the collaboration for a longer period prior to its release to the public.

- Purchase, without delay, the needed hardware servers to support the DR3 development. This is important to allow sufficient time to install, test, and commission the hardware and load DR3. As there are few technical concerns about the hardware plan, this should be not delayed for completion of the up-dated requirements and implementation plan.

- Add data validation tests to the implementation plan, including an estimate of manpower needed to perform those tests.

2.3 Is the hardware implementation plan reasonable? In particular, is the design of hot spares, the types of RAID arrays, etc appropriate for meeting the requirements? Is the plan for scaling the CAS to meet the needs at the end of the survey sensible? Are the hardware costs well understood?

Findings

- *The hardware at FNAL to support the CAS has been insufficient.* Delays in the data releases, at least to the collaboration, have been contributed to by not having the necessary hardware in hand.

- *Security and support issues have hampered siting the primary CAS at FNAL.* While there was a desire to have the primary DR1 server based at FNAL, that in fact did not happen partly due to the failure to understand support issues at FNAL. Problems dealing with FNAL security requirements led to a delay of many weeks. Note that just such a possible problem concerning FNAL security was raised in the 2000 review report: "We also note that restricted access implies an access control mechanism (passwords, data encryption, etc.) and that access control issues may become complex if Fermilab hosts the public archive and is required to implement enhanced computer security measures."

- *The current hardware plan is complicated due to fiscal constraints.* The ideal plan, as presented, would involve more machines, with each machine playing fewer roles, than the adopted plan.

Comments

- *We commend the creative hardware plan presented given the fiscal constraints.* The adoption of multiple roles for each given machine, and the expectation that larger disks will be available for the final data release, seem to be reasonable responses to a limited budget.

- *We have no technical concerns.* The choice of CPUs, raid arrays, etc. appears to be appropriate. There is a certain sense of "fragility" — the sense that everything must work just so for the schedule to hang together — however the plan appears to be workable.

- *The hardware impact of some new features have not been considered.* For example, their was no discussion of the disk space requirements of CasJobs/MyDB, and its not clear that feature was even considered for the FNAL hardware plan. Again, this points to the lack of performance requirements, and their consideration in drafting the hardware plan.

- *The plan is not sufficiently detailed to fully evaluate.* While a detailed list of hardware was presented, many important issues, such as support (short and long-term), integration with FNAL, etc. were not presented, and indeed have not apparently been thought through, in sufficient detail. Further, again the lack of performance requirements makes it impossible to determine whether the hardware plan can meet the performance needs for both the public and the collaboration.

Recommendations

- Revisit the hardware plan, with an eye toward manpower, support, requirements, integration, security, collaboration needs, etc. All aspects of basing the primary CAS server at FNAL must be worked through in full detail if experiences like the first attempt of basing DR1 at FNAL are to be avoided.

- Evaluate the possible role for mass storage. Having mass storage based backups rather than fully operational backup machines may allow for a better use of the limited budget.

- Understand the requirements for a proper testbed at JHU.

2.4 Are the resources, including manpower at JHU and Fermilab, to deploy and operate the CAS reasonable? Is the level of interaction between JHU and Fermilab sufficient?

Findings

- *Both JHU and FNAL personnel have been doing heroic jobs.* It was obvious from the Review presentations that an enormous amount of work has been done with very few resources. JHU has developed extremely fast access mechanisms for such a large database, as well as very sophisticated user query tools. They have also done an excellent job in developing the CAS data filtering and loading tools. FNAL has been working very hard to load the CAS, but have been hampered by insufficient disk space and upstream file problems. They have developed very creative plans for utilizing their hardware to handle new data releases.

- *FNAL does not have adequate manpower to support the CAS operations.* FNAL does not have sufficient resources to support the CAS as the primary site. Development resources from JHU are being called upon repeatedly to configure and administer the database. To meet the 24/7 uptime requirements will necessitate having sufficient staff to establish a reliable configuration with adequate on-call support.

Comments

- *Understanding the requirements is necessary to understand the manpower needs at JHU.* JHU development is supported by direct ARC funding and also by other independently funded grants. It will be necessary to separate the requirements for the CAS from the requirements for the user interface to be able to analyze whether there are sufficient JHU resources to support completing required CAS functionality in time to support the milestone releases.

- *FNAL and JHU have failed to agree on a prioritized list of tasks.* As mentioned in earlier sections of this report, there are inadequate current operational and performance requirements and no implementation schedule, therefore work is proceeding without an agreed upon prioritization.

- *We commend the collaboration with CSS/CSI.* The FNAL/CD brings much needed expertise, and support infrastructure, for the deployment of a fully-supported public data set. The SDSS management clearly recognizes the value of such a collaboration.

- *No plan for the long-term (beyond DR5) support of the CAS was presented.* The siting, staffing, and hardware for the CAS beyond the life of the survey

needs to be planned for now. Long-term support may well constrain the choice of architecture, security issues, etc, and these need to be understood as the system in being designed.

Recommendations

- Improve communication between JHU and FNAL concerning priorities: updated requirements and an implementation plan will facilitate this. Specific functionality deliverables should be tied to milestones on the schedule. Frequent status reviews should identify risks to the schedule and explain risk mitigation plans.
- Supply one full-time support person at FNAL for the CAS with Windows and SQL Server expertise. This expertise will be critical to smoothing the handoffs of new deliveries from JHU, and for resolving database problems at the primary site.
- Find a full-time Head of Data Distribution. The task of developing and monitoring a feasible schedule for the deployment of the remaining data releases will be crucial to the success of FNAL as the primary site. This in no way is meant to denigrate the job being done by the current acting head, who has in fact considerably improved communication between JHU and FNAL. It is merely a statement that we view the position as a full-time job, at least for the life of the survey.
- Seek further collaboration with FNAL/CD. The CD has valuable experience with Windows systems, fault tolerant system architectures, and with lights-out operations that will be important to a successful CAS primary site deployment.
- Write MOUs between EAG/JHU, and between EAG/CSS, so that the support roles for all groups are clearly understood, both in the near and long term.
- Plan now for the long-term support of the CAS at FNAL beyond DR5. Final system architecture and staff resources will have to be investigated.